arXiv:1603.01551v1 [stat.CO] 4 Mar 2016

Perfect and ε-Perfect Simulation Methods for the One Dimensional Kac Equation

J. N. Corcoran, D. Jennings, and P. VaughanMiller
University of Colorado *

March 7, 2016

Abstract

We review the derivation of the Kac master equation model for random collisions of particles, its relationship to the Poisson process, and existing algorithms for simulating values from the marginal distribution of velocity for a single particle at any given time. We describe and implement a new algorithm that efficiently and more fully leverages properties of the Poisson process, show that it performs at least as well as existing methods, and give empirical evidence that it may perform better at capturing the tails of the single particle velocity distribution. Finally, we derive and implement a novel "ε-perfect sampling" algorithm for the limiting marginal distribution as time goes to infinity. In this case the importance is a proof of concept that has the potential to be expanded to more interesting (DSMC) direct simulation Monte Carlo applications.

1 Introduction

In [7], Kac introduced an N particle stochastic process model for a spatially homogeneous gas that evolves according to a specific collision process. He proved that, as $N \to \infty$, the marginal probability density function (pdf) for a single particle at time t approaches the solution of a simplified version of the famed Boltzmann equation. The N-dimensional probability density function for all particles is described by the **Kac master equation** which can be derived as a random walk with transitions occurring at times of a Poisson process. In this paper we take advantage of properties of the Poisson process, including process "thinning" and the conditional uniformity of event times to efficiently simulate values from the velocity density of any single particle at any given time $t > 0$. We compare our results to those from the familiar Nanbu, Nanbu-Babovsky, and DSMC (direct simulation Monte Carlo) algorithms. Additionally, we introduce an "ε-perfect sampling algorithm" that uses "coupling from the past" techniques to simulate or sample from the velocity distribution at

*Postal Address: Department of Applied Mathematics, University of Colorado, Box 526 Boulder CO 80309-0526, USA; email: corcoran@colorado.edu, phone: 303-492-0685

[0]Keywords: MCMC, Kac Model, rarefied gas flows, perfect simulation, DSMC
AMS Subject classification: 60J10, 65C05, 76P05

time $t = \infty$ up to a user defined accuracy ε that can be taken so small as to be essentially "perfect" within the accuracy available by a given machine precision.

In Section 2, we review the derivation of the N particle Kac master equation and show its relationship to the Poisson process. In Section 3 we very briefly review the probability density function of any one given particle as a marginal density of the joint probability density function solution to the Kac master equation. In Section 4 we review three existing simulation methods for the marginal density, introduce the exact Poisson method, and make an empirical comparison between all methods. In Section 5, we briefly review the idea of perfect simulation (perfect sampling) and describe a novel "ε-perfect" algorithm that applies to the Kac model.

2 The Kac Master Equation

Consider N one-dimensional particles, each with mass m, with velocities $v_1, v_2, \ldots, v_N$. Define $\vec{v} = (v_1, v_2, \ldots, v_N)$. The total kinetic energy of the ensemble of particles is

$$E = \frac{m}{2} \sum_{i=1}^{N} v_i^2.$$

For simplicity, we assume that $m = 2$. Thus, we have $E = \sum_{i=1}^{N} v_i^2$ or, equivalently, $|\vec{v}| = \sqrt{E}$.

If we assume that kinetic energy is conserved as these velocities evolve through particle collisions, points $\vec{v}$ with always be found on the sphere $S^{N-1}(\sqrt{E})$. We now describe a particular evolution of velocities, proposed by Kac [7] in two stages as in [3] The first stage is to describe as a particular discrete time random walk known as the "Kac walk" and the second stage is to modify the random walk so that it is taking place in continuous time.

2.1 The Kac Walk

The Kac Walk:

1. Randomly select a pair of indices (i, j) from $\{1, 2, \ldots, N\}$. There are $\binom{N}{2}$ ways to do this, so a particular pair (i, j) is selected with probability $\left[\binom{N}{2}\right]^{-1}$
2. Randomly select a *scattering angle* θ from a density $\rho(\theta)$ on $[0, 2\pi)$. (Often θ is assumed to be uniformly distributed on $[0, 2\pi)$. For an upcoming derivation, it will be convenient to define $\rho(\theta)$ on $[-\pi, \pi)$ and to assume that it is symmetric about 0.)
3. Update the velocities for particles i and j through the rotation

   $$(v_i, v_j) \to (v_i', v_j') := (v_i \cos\theta + v_j \sin\theta, -v_i \sin\theta + v_j \cos\theta).$$

   Return to Step 1.

It is easy to see that total kinetic energy is conserved by the Kac walk (though momentum is not) and therefore that these three steps create a random walk on $S^{N-1}(\sqrt{E})$. We will now derive the probability density function, f_k for the vector of N velocities after k "collisions" or iterations of the Kac walk.

Let $g : S^{N-1}(\sqrt{E}) \to \mathbb{R}$ be any continuous function. Define an operator Q_N on such functions as the conditional expectation

$$Q_N g(\vec{v}) := \mathsf{E}[g(\vec{V}_{k+1})|\vec{V}_k = \vec{v}]$$

where $\vec{V}_k$ is the random vector of velocities at the kth iteration of the Kac walk.

Conditioning on the selected angle and particles, this becomes

$$\begin{aligned}
Q_N g(\vec{v}) &= \mathsf{E}[g(\vec{V}_{k+1})|\vec{V}_k = \vec{v}] \\
&= \int_{-\pi}^{\pi} \sum_{i<j} \mathsf{E}\left[g(\vec{V}_{k+1}) \middle| \vec{V}_k = \vec{v}, \begin{array}{c}\text{particles } i \text{ and } j \\ \text{are selected}\end{array}, \begin{array}{c}\text{angle } \theta \text{ is} \\ \text{selected}\end{array} \right] \cdot \left[\begin{pmatrix} N \\ 2 \end{pmatrix} \right]^{-1} \rho(\theta)\, d\theta \\
&= \int_{-\pi}^{\pi} \sum_{i<j} g(v_1, \ldots, v_{i-1}, v_i', v_{i+1}, \ldots, v_{j-1}, v_j', v_{j+1}, \ldots, v_N) \cdot \left[\begin{pmatrix} N \\ 2 \end{pmatrix} \right]^{-1} \rho(\theta)\, d\theta \\
&= \int_{-\pi}^{\pi} \sum_{i<j} g(R_{ij}(\theta)\, \vec{v}) \cdot \left[\begin{pmatrix} N \\ 2 \end{pmatrix} \right]^{-1} \rho(\theta)\, d\theta
\end{aligned}$$

where $R_{ij}(\theta)$ is the $N \times N$ rotation matrix that is the identity matrix altered to have $\cos\theta$ as the (i,i) entry, $\sin\theta$ as the (i,j) entry, $-\sin\theta$ as the (j,i) entry, and $\cos\theta$ as the (j,j) entry.

Let f_0 be an initial probability density function for the N velocities and let f_1 be the density after one iteration. Define random vectors $\vec{V}_0 \sim f_0$ and $\vec{V}_1 \sim f_1$. Note that

$$\mathsf{E}[g(\vec{V}_1)] = \int_S g(\vec{v})\, f_1(\vec{v})\, d\vec{v},$$

where S is shorthand for $S^{N-1}(\sqrt{E})$. Also note that

$$\begin{aligned}
\mathsf{E}[g(\vec{V}_1)] &= \mathsf{E}[\mathsf{E}[g(\vec{V}_1)|\vec{V}_0]] \\
&= \textstyle\int_S \mathsf{E}[g(\vec{V}_1)|\vec{V}_0 = \vec{v}] \cdot f_0(\vec{v})\, d\vec{v} \\
&= \textstyle\int_S Q_N g(\vec{v}) f_0(\vec{v})\, d\vec{v}
\end{aligned}$$

Thus, we have that

$$\int_S g(\vec{v})\, f_1(\vec{v})\, d\vec{v} = \int_S Q_N g(\vec{v}) f_0(\vec{v})\, d\vec{v}. \tag{1}$$

It is straightforward to verify that Q_N is a self-adjoint operator in the sense that

$$\int_S Q_N g(\vec{v}) f(\vec{v})\, d\vec{v} = \int_S g(\vec{v}) Q_N(\vec{v}) f(\vec{v})\, d\vec{v} \tag{2}$$

for all continuous g and for all N dimensional densities f. Therefore, combining (1) and (2), we have

$$\int_S g(\vec{v}) f_1(\vec{v})\, d\vec{v} = \int_S g(\vec{v})\, Q_N f_0(\vec{v})\, d\vec{v}. \tag{3}$$

Since (3) holds for all $g : S^{N-1}(\sqrt{E}) \to \mathbb{R}$, we have that $f_1 = Q_N f_0$. Similarly, the probability density function for the velocities after k collisions is given by

$$f_k = Q_N^k f_0. \tag{4}$$

2.2 Time and the Poisson Process

As of now, we have a pdf for the N-dimensional vector of velocities evolving at the discrete times of a random walk. We now will allow the velocities to evolve as an N-dimensional continuous time Markov process. The pdf will now be denoted by $f(\vec{v}, t)$.

Kac proposed in [7] that collisions happen in the N-particle ensemble at times of a Poisson process. Specifically, it is assumed that the expected number of particles that collide in a vanishingly small time step Δt is $\lambda N \Delta t$ for some $\lambda > 0$. This means that the expected number of collisions between pairs of particles is $\lambda N \Delta t/2$ and so the assumption is that collisions happen according to a Poisson process with rate $\lambda N/2$. Among other things, this implies that, as $\Delta t \searrow 0$, the probability of exactly one collision in any time interval $(t, t + \Delta t)$ is $\lambda N \Delta t/2 + o(\Delta t)$, the probability of no collisions is $1 - \lambda N \Delta t/2 + o(\Delta t)$, and the probability of two or more is $o(\Delta t)$. (Here, $o(\Delta t)$ represents a vanishing function $g(\Delta t)$ such that $g(\Delta t)/\Delta t \to 0$ as $\Delta t \searrow 0$.)

Given the velocity density $f(\vec{v}, t)$, at time t, the density at time $t + \Delta t$ as $\Delta t \searrow 0$ is, informally, the mixture density

$$\begin{aligned} f(\vec{v}, t + \Delta t) \;=\; & p \cdot (\text{ previous density after having gone through a collision }) \\ & +(1-p) \cdot (\text{ previous density }) \end{aligned}$$

where p is the probability of a collision in the ensemble during the time interval $(t, t + \Delta t)$ as $\Delta t \searrow 0$.

Formally, this is written as

$$\begin{aligned} f(\vec{v}, t + \Delta t) &= (\lambda N \Delta t/2 + o(\Delta t)) \cdot Q_N f(\vec{v}, t) + (1 - \lambda N \Delta t/2 + o(\Delta t)) \cdot f(\vec{v}, t) \\ &= \tfrac{\lambda N \Delta t}{2}\, Q_N f(\vec{v}, t) + f(\vec{v}, t) - \tfrac{\lambda N \Delta t}{2}\, f(\vec{v}, t) + o(\Delta t) \end{aligned}$$

which implies that

$$\frac{f(\vec{v}, t + \Delta t) - f(\vec{v}, t)}{\Delta t} = \frac{\lambda N [Q_N - I]}{2} f(\vec{v}, t) + \frac{o(\Delta t)}{\Delta t}$$

where I is the identity operator.

Letting $\Delta t \searrow 0$, $\beta = \lambda/2$, and specifying an initial condition, we get the **Kac Master Equation**:

$$\begin{aligned} \tfrac{d}{dt} f(\vec{v}, t) &= \beta N [Q_N - I] f(\vec{v}, t) \\ f(\vec{v}, 0) &= f_0(\vec{v}) \end{aligned}$$

which has solution

$$f(\vec{v}, t) = e^{\beta N [Q_N - I] t} f_0(\vec{v}) = e^{-\beta N t} \sum_{i=0}^{\infty} \frac{(\beta N t)^i}{i!} Q_N^i f_0(\vec{v}).$$

3 The One Dimensional Kac Equation

We assume henceforth that the scattering angle density, $\rho(\theta)$, is the uniform density on the interval $(0, 2\pi)$. We write $\theta \sim \text{unif}(0, 2\pi)$. Consider the marginal density for particle 1,

$$f_N(v_1, t) = \int_{S^{N-2}(\sqrt{E - v_1^2})} f(\vec{v}, t)\, d\vec{v}_{-1} \tag{5}$$

where $\vec{v}_{-1} := (v_2, v_3, \ldots, v_N)$. We will make the typical assumption that $E = N$ since the total kinetic energy is proportional to the number of particles in the system.

Kac [7] showed, under some assumptions[1] on the initial density $f(v, 0)$, that this single particle marginal density converges, as $N \to \infty$, to the solution of a one-dimensional spatially homogeneous Boltzmann-type equation:

$$\frac{df}{dt}(v, t) = \frac{\lambda}{2\pi} \int_{-\infty}^{\infty} \int_0^{2\pi} [f(v', t) f(w', t) - f(v, t) f(w, t)]\, d\theta\, dw, \tag{6}$$

where

$$v' = v \cos\theta + w \sin\theta \qquad \text{and} \qquad w' = -v \sin\theta + w \cos\theta.$$

Equation (6) is known as the **Kac equation** or the **Kac-Boltzmann equation**.

Note that (6) may be rewritten as

$$\frac{\partial f}{\partial t} = \frac{1}{\varepsilon} [P(f, f) - f] \tag{7}$$

for $\varepsilon = 1/\lambda$, and $P(f, f)$ defined as the *bilinear collision operator*

$$P(f, f) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_0^{2\pi} f(v', t) f(w', t)\, d\theta\, dw.$$

[1] Let ϕ be a one dimensional density. Informally, an n-dimensional density ϕ_n, symmetric in its arguments, is "ϕ-chaotic" if, as $n \to \infty$ any k-dimensional marginal of ϕ_n acts as a product of k copies of ϕ. In [7], Kac proved that if $f^{(N)} = f(\vec{v}, t)$ solves the Kac Master equation, then $f^{(N)}$ is f-chaotic where f solves the Kac-Boltzmann equation. Thus, we can consider any univariate marginal of $f(\vec{v}, t)$ as $N \to \infty$ as a solution to (6). For a more formal explanation, please see [4].

If the initial distribution $f_0(v) := f(v,0)$ on $(-\infty,\infty)$ is taken to be $f_0(v) = \frac{2}{\sqrt{\pi}} v^2 e^{-v^2}$, and if $\lambda = \sqrt{\pi}/2$, the solution to (6) is known [8] to be

$$f(v,t) = \frac{1}{\sqrt{\pi}} \left[\frac{3}{2}(1 - C(t))\sqrt{C(t)} + (3C(t) - 1)(C(t))^{3/2} v^2 \right] e^{-C(t)v^2} \tag{8}$$

where

$$C(t) = [3 - 2\exp(-\sqrt{\pi}t/16)]^{-1}.$$

It is easy to see that

$$\lim_{t\to\infty} f(v,t) = \frac{1}{\sqrt{3\pi}}\, e^{-\frac{1}{3}v^2} \qquad \text{for } -\infty < v < \infty. \tag{9}$$

We will use (8) and (9) to test several Monte Carlo algorithms.

4 Simulation Methods for $f(v,t)$

In Sections 4.1 and 4.2 we review existing methods for drawing (simulating) velocity values from the single particle marginal distribution $f_N(v,t)$ for finite t. In Section 4.3 we propose a new method that takes advantage of properties of the Poisson collision process. Our goal here is to have a method efficient enough to be able to push N quite large so that we may see approximately $f(v,t)$.

4.1 The Nanbu and Nanbu-Babovsky Algorithms

In [9] Nanbu proposes an algorithm to draw from $f_N(v,t)$ based on the probabilistic interpretation of colliding particles and on the discretization

$$f_{n+1} = (1 - \lambda\Delta t)\, f_n + \lambda\Delta t P(f_n, f_n) \tag{10}$$

of (7). Here, $f_n = f_n(v) = f(v,t_n)$, $t_0 = 0$, and $t_n = t_0 + n\Delta t$ for some small increment $\Delta t > 0$. The interpretation of (10) is that, if Δt is small enough so that $0 \le \lambda\Delta t \le 1$, "particle 1" (or any fixed given particle)

- fails to undergo a collision and does not have any velocity update with probability $1 - \lambda\Delta t$, or
- experiences a collision with probability $\lambda\Delta t$ and has a new velocity generated according to the collision process described by $P(f,f)$.

For the following algorithm descriptions, we assume that Δt divides into t evenly.

Nanbu's Algorithm

For $i = 1, 2, \ldots, N$, let $v_i^{(n)}$ denote the velocity for particle i at time step $n = 0, 1, 2, \ldots$.

1. Sample initial velocities, $v_1^{(0)}, v_2^{(0)}, \ldots, v_N^{(0)}$, for all N particles independent and identically distributed (iid) from $f_0(v)$. Set $s = 0$.

2. For $i = 1, 2, \ldots, N$,

   - With probability $1 - \lambda\Delta t$ do nothing to the velocity of the ith particle. Set $v_i^{(n+1)} = v_i^{(n)}$.
   - With probability $\lambda\Delta t$,
     - Randomly select a second particle from among the remaining particles in $\{1, 2, \ldots, N\}$.
     - Draw an angle value $\theta \sim \rho(\theta)$.
     - Set $v_i^{(n+1)} = v_i^{(n)} \cos\theta + v_j^{(n)} \sin\theta$

3. Set $s = s + \Delta t$ and $n = n + 1$. If $s < t$, return to Step 2.

As kinetic energy is not conserved in Nanbu's Algorithm, Babovsky [1] modified the algorithm to update both v_i and v_j as particles i and j collide. Additionally, at each time step, the number of collisions processed is taken to be equal to the expected number of collision pairs in an interval of length Δt. As the expected number of collisions is $\lambda N \Delta t$, the expected number of collision pairs is $\lambda N \Delta t/2$. This is then probabilistically rounded using the following function.

$$Round[x] = \begin{cases} \lfloor x \rfloor & , \text{ with probability } \lfloor x \rfloor + 1 - x \\ \lfloor x \rfloor + 1 & , \text{ with probability } x - \lfloor x \rfloor \end{cases}$$

It is easy to see that, for any given x, the expected value of $Round[x]$ is x.

Nanbu-Babovsky Algorithm

For $i = 1, 2, \ldots, N$, let $v_i^{(n)}$ denote the velocity for particle i at time step $n = 0, 1, 2, \ldots$.

1. Sample initial velocities, $v_1^{(0)}, v_2^{(0)}, \ldots, v_N^{(0)}$, for all N particles independent and identically distributed (iid) from $f_0(v)$. Set $s = 0$.

2. For $i = 1, 2, \ldots, N$,

   - Select $m = Round[\lambda N \Delta t/2]$.
   - Randomly select m pairs, (i, j) of particle indices, without replacement, from among all possible pairs.
   - For each selected pair (i, j), draw an angle value $\theta \sim \rho(\theta)$ and set

$$\begin{aligned} v_i^{(n+1)} &= v_i^{(n)} \cos\theta + v_j^{(n)} \sin\theta \\ v_j^{(n+1)} &= -v_i^{(n)} \sin\theta + v_j^{(n)} \cos\theta. \end{aligned}$$

- For the remaining unselected indices k, set $v_k^{(n+1)} = v_k^{(n)}$.

3. Set $s = s + \Delta t$ and $n = n + 1$. If $s < t$, return to Step 2.

Both the Nanbu and the Nanbu-Babovsky algorithms can be shown to converge to the solution of the discretized Kac equation given by (10). Nanbu's algorithms allows each particle to potentially undergo a collision only once per time step and the Nanbu-Babovsky scheme allows each pair of particles to potentially undergo a collision once per time step.

4.2 Bird's DSMC Algorithm

The Direct Simulation Monte Carlo (DSMC) approach due to Bird [2] was first proposed in the late 60's and has been shown [10] to converge to the solution of (7). It differs from the Nanbu-Babovsky algorithm in that any particular collision pair may be selected multiple times in the time interval of length Δt. Since there are $\lambda N \Delta t/2$ collisions expected in that interval, the average time between collisions is taken to be $\Delta t_c := \Delta t/[\lambda N \Delta t/2] = 2/(\lambda N)$.

Bird's DSMC Algorithm

1. Sample initial velocities, $v_1^{(0)}, v_2^{(0)}, \ldots, v_N^{(0)}$, for all N particles independent and identically distributed (iid) from $f_0(v)$. Set $s = 0$.

2. Process one collision as follows.

   - Select a pair (i, j) of particles at random from among all possible pairs.
   - For each selected pair (i, j), draw an angle value $\theta \sim \rho(\theta)$ and set

$$\begin{aligned} v_i^{(n+1)} &= v_i^{(n)} \cos\theta + v_j^{(n)} \sin\theta \\ v_j^{(n+1)} &= -v_i^{(n)} \sin\theta + v_j^{(n)} \cos\theta. \end{aligned}$$

   - For the remaining unselected indices k, set $v_k^{(n+1)} = v_k^{(n)}$.

3. Set $s = s + \Delta t_c$ and $n = n + 1$. If $s < t$, return to Step 2.

Bird's DSMC techniques comprise a much wider class of extremely useful and popular algorithms than the one detailed here. As it stands, the one given here is essentially the Nanbu-Babovsky algorithm where pairs are selected with replacement.

4.3 Exact Simulation of $f_N(v, t)$ Based on Properties of the Poisson process

In this section, we simulate exactly from $f_N(v, t)$ using the actual Poisson numbers of collisions between important time points as opposed to the expected numbers of collisions. We do not partition the interval $(0, t)$ as in the previously described algorithms but instead focus on collision times for the fixed "particle 1".

Recall that the N-particle ensemble is assumed to undergo collisions at times of a Poisson process with rate λN. As each collision involves 2 particles, this means that pairs of particles are colliding according to a Poisson process with rate $\lambda N/2$. The probability that particle 1 is involved in any one of these collisions is $(N-1)/\binom{N}{2} = 2/N$. So, the particle 1 collision process is a thinned Poisson process with rate $(\lambda N/2) \cdot 2/N = \lambda$. Therefore, we begin by simulating a Poisson rate λt number of collisions for particle 1 in the time interval $(0, t)$.

Though intercollision times for particle 1 are exponential with rate λt, once the number of collisions during $(0, t)$ is fixed, the times of those collisions are uniformly distributed over the interval $(0, t)$. So, rather than divide up the interval $(0, t)$ into Δt increments, we consider only the times when particle 1 is involved in a collision by first simulating $K \sim Poisson(\lambda t)$, then simulating independent $U_1, U_2, \ldots, U_K$ from the uniform distribution over $(0, t)$ and finally considering the ordered time points $T_1, T_2, \ldots, T_K$ corresponding to the ordered U_i. (For example, $T_1 = \min(U_1, U_2, \ldots, U_K)$ and $T_K = \max(U_1, U_2, \ldots, U_K)$.)

To update the velocity for particle 1 at time T_1, we must choose, at random, a particle from the remaining $N - 1$ particles for collision. As this "sub-ensemble" of particles has been evolving over the time interval $(0, T_1)$, we will process $K_1 \sim Poisson(\lambda(N-1)T_1/2)$ collisions involving only these particles without any regard for the times of these collisions and then select one at random for collision with particle 1. Continuing in this manner, we get the following algorithm.

Exact Poisson Algorithm

1. Sample initial velocities, $v_1^{(0)}, v_2^{(0)}, \ldots, v_N^{(0)}$, for all N particles independent and identically distributed (iid) from $f_0(v)$.

2. Simulate $K \sim Poisson(\lambda t)$. This is the number of collisions involving particle 1 during $[0, t]$. Suppose that $K = k$.

3. Simulate $U_1, U_2, \ldots, U_k \overset{iid}{\sim} \text{unif}(0, t)$ collision times for particle 1. Denote the ordered $U_1, U_2, \ldots, U_k$ as $T_1, T_2, \ldots, T_k$ where $T_1 = \min(U_1, U_2, \ldots, U_k)$ and $T_k = \max(U_1, U_2, \ldots, U_k)$.

4. Let $T_0 = 0$.

   For $i = 1, 2, \ldots, k$,

   - Simulate $K_i \sim Poisson(\lambda(N-1)(T_i - T_{i-1})/2)$. This is the number of "non particle 1" collisions happening during the time interval $[T_{i-1}, T_i]$. Suppose that $K_i = k_i$.
   - Process k_i collisions for the $(N-1)$ particle ensemble as follows.
     For $j = 1, 2, \ldots, k_i$
     - Select a pair (r, s) of particles at random from among all possible pairs or particles with labels in $\{2, 3, \ldots, N\}$.
     - For each selected pair (r, s), draw an angle value $\theta \sim \rho(\theta)$ and set new velocities as

$$\begin{aligned} v_r &= v_r \cos\theta + v_s \sin\theta \\ v_s &= -v_r \sin\theta + v_s \cos\theta \end{aligned}$$

where v_r and v_s on the right-hand side of these equations are the pre-collision velocities for particles r and s.

- Randomly select one particle from $\{2, 3, \ldots, N\}$ and process a collision with particle 1 by assigning new velocities

$$\begin{aligned} v_r &= v_r \cos\theta + v_s \sin\theta \\ v_s &= -v_r \sin\theta + v_s \cos\theta \end{aligned}$$

where r is the selected index in $\{2, 3, \ldots, N\}$ and v_1 and v_r on the right-hand side of these equations are the pre-collision velocities for particles 1 and r.

Note that, unlike the Nanbu, Nanbu-Babovsky, and Bird algorithms, we are not using any sort of time counter or partitioning of the time interval $[0, t]$. Furthermore, we are actually processing a Poisson number of collisions, as assumed by the derivation of the Kac equation, in any given time interval rather than simply the expected number of collisions under the Poisson assumption as in Bird's algorithm. Finally, we note that we are saving on collisions by not propagating the system in the time interval $[T_K, t]$ as we are trying to sample the velocity of particle 1 at time t and this is unchanged in this interval. This gives an expected savings of $(N/2)(1 - e^{-\lambda t})$ collisions per simulation repetition.

4.4 Simulation Results

In order to assess the performance of the different algorithms, we simulated repeated draws from the velocity distribution of the first particle, v_1, for various values of N, and compared the results to the exact solution (for $N \to \infty$) given by (8). Specifically, for each algorithm we simulated N iid velocities and propagated the ensemble forward according to the algorithms described in Sections 4.1, 4.2, and 4.3, collecting the ultimate values for v_1 at various times t. In all cases, we were able to achieve reasonable accuracy for a surprisingly small number of particles. However, since the accuracy of the Nanbu and Nanbu-Babovsky algorithms both depend on how finely the time interval is partitioned and since the Nanbu-Babovsky sampling with replacement step is relatively inefficient, we will focus our comparisons mainly on the Nanbu, Bird, and Poisson algorithms.

We will express the accuracy of our numerical results using an estimate of the **total variation norm** (TVN) distance between the true curve and the estimated distributions. The total variation norm distance between two probability measures π_1 and π_2 is defined to be

$$||\pi_1 - \pi_2|| = \frac{1}{2} \sup_{A \in \mathcal{B}} |\pi_1(A) - \pi_2(A)|,$$

where $\mathcal{B}$ is the set of Borel sets in the appropriate space and $\pi_i(A) = \int_A \pi_i(x)\, dx$. If π_1 and π_2 are discrete measures, then the total variation norm can be written simply as

$$||\pi_1 - \pi_2|| = \frac{1}{2} \sum_x |\pi_1(x) - \pi_2(x)|. \tag{11}$$

In this paper, we discretized both the true and simulated densities by a natural histogram binning procedure in order to use (11) to give an estimated TVN distance between the

Figure 1: Histograms of 100,000 values of v_1 at $t = 2$ with $N = 5$ for $\Delta t = 0.01$ (left) and $\Delta t = 1.0$ (right). Curve is exact solution for $N = \infty$ given by (8).

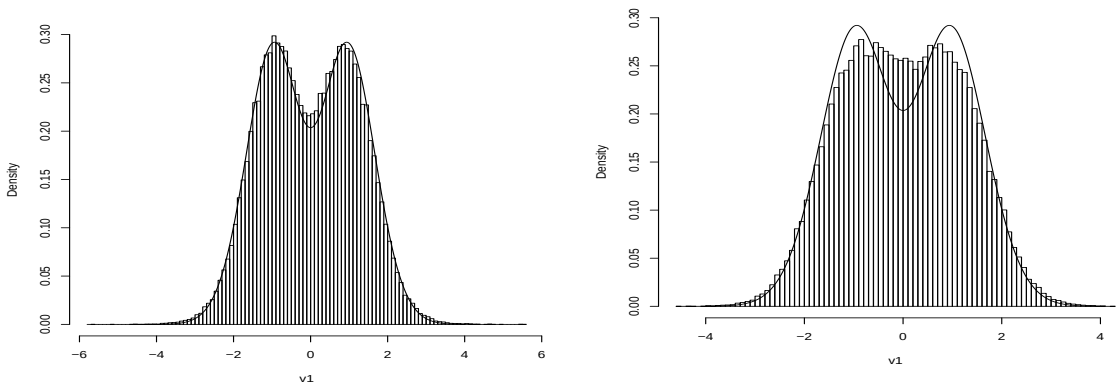

target density and our simulation results. As this estimate depends on the bin widths used, the actual values of our estimated TVN are not as interesting as the relative values. We will keep our bin widths constant throughout.

For example, for $t = 2$, using Nanbu's algorithm with $N = 5$ produced distributions for v_1 shown in Figure 1. In both cases the "target density" given by (8) is shown as a solid curve. It is no surprise that the smaller simulation time step gives better results, but to illustrate the TVN error, the estimated values of TVN for the left and right graphs are 0.0187 and 0.0482, respectively. Figure 2 shows estimated TVN, for several values of N, as the simulation time step size decreases. Each point on the graph was created by averaging 100 values of TVN, each having been computed from a sample of size $100,000$. The height of the horizontal line just below 0.01 shows the average TVN for 100 samples of size $100,000$ for Bird's DSMC algorithm (which doesn't depend on Δt) for $N = 50$. This was indistinguishable to several decimal places from similarly estimated TVN for larger N, for the Poisson algorithm for $N \geq 50$, and for a baseline estimated TVN for samples of size $100,000$ from the standard normal distribution. (Recall that we do not expect the estimated TVN to necessarily approach zero even in the case of exact simulation from the normal distribution due to the discretization used in its calculation.) Figure 3 shows the average TVN values (again averaging 100 values each computed from a sample of size $100,000$) for the three algorithms as a function of the number of particles in the simulations. Here we can see that the Poisson algorithm consistently performs better than both the Nanbu and Bird algorithms for small N. That is, less particles are needed in order to approximate (8) when using the Poisson algorithm. Distributions for v_1, estimated by the DSMC and Poisson algorithms are shown in Figure 4. While they are similar both visually and by the TVN measure, the Poisson algorithm appears to be performing slightly better in the tails, as shown in Figure 5 for the case of the upper tails.

Figure 2: Estimated TVN for Nanbu's Algorithm as a function of k for $\Delta t = 2/2^k$

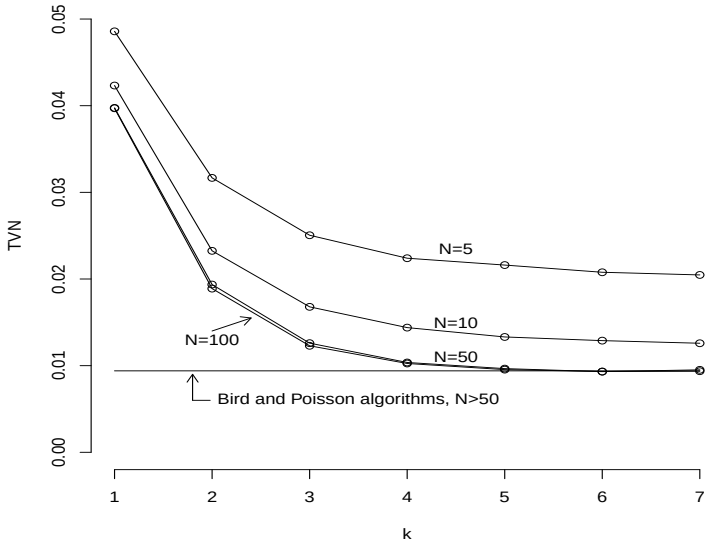

Figure 3: Estimated TVN for the Nanbu, Bird, and Poisson algorithms as a function of N

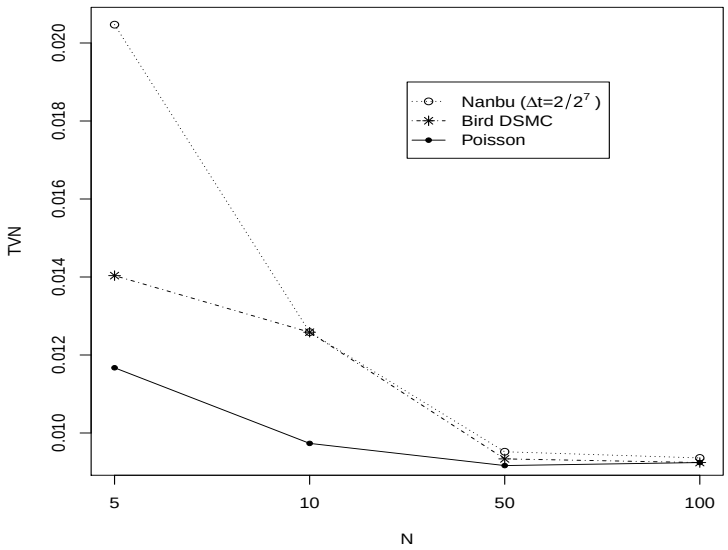

Figure 4: Histograms of 100,000 Values of v_1 at $t = 2$ with $N = 50$ for the DSMC algorithm (left) and the Poisson algorithm (right).

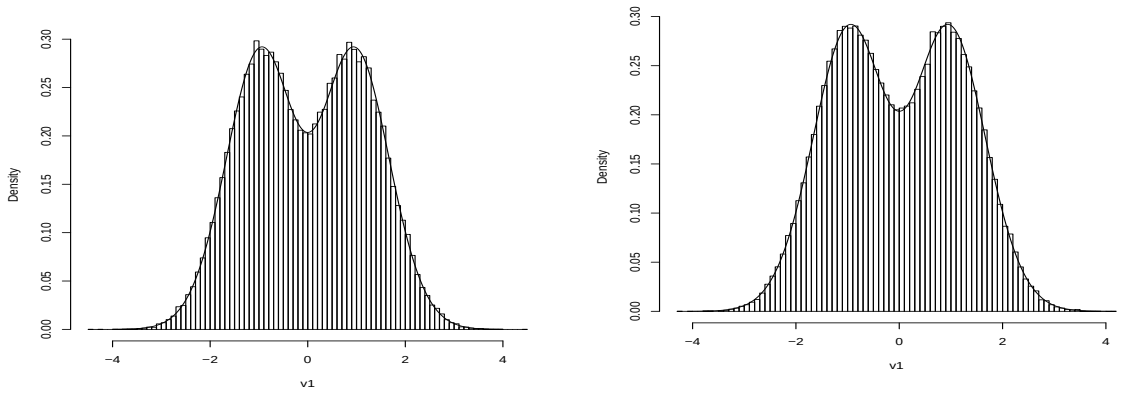

Figure 5: Histograms of 100,000 upper tail values of v_1 at $t = 2$ with $N = 1000$ using Bird's DSMC algorithm (left) and the Poisson algorithm (right).

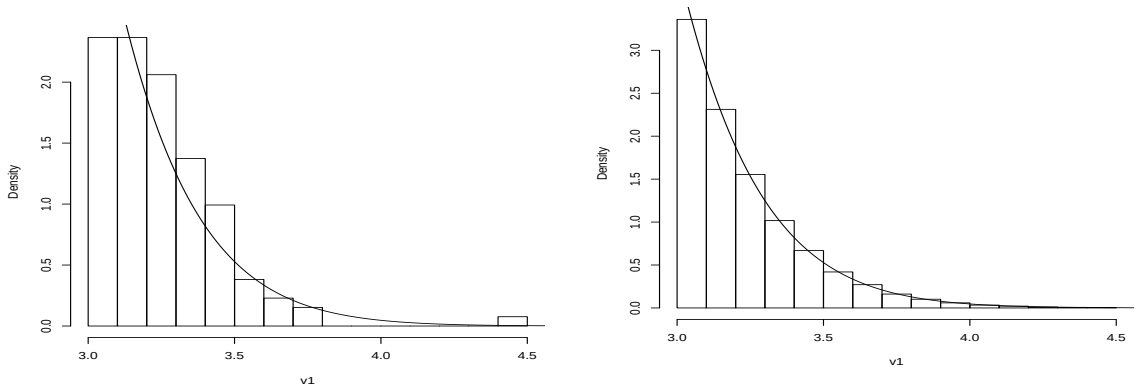

5 An ε-Perfect Simulation of $f(v, \infty)$

"Perfect sampling" (or "perfect simulation") is the name applied to a class of Markov chain Monte Carlo (MCMC) algorithms which enable one to draw values exactly from the stationary (limiting) distribution π of a Markov chain. This is in contrast to the more typical MCMC algorithms in which transitions are simulated for "a long time" after which output is collected giving only approximate draws from π.

5.1 Perfect Simulation

The essential idea of most of these approaches is to find a random epoch $-T$ in the past such that, if one constructs sample paths (according to a transition density $p(x, y)$ that has stationary/limiting distribution π) from every point in the state space starting at $-T$, all paths will have met or "coupled" by time zero. The common value of the paths at time zero is a draw from π. Intuitively, it is clear why this result holds with such a random time T. Consider a chain starting anywhere in the state space at time $-\infty$. At time $-T$ it must pick *some* value x, and from then on it follows the trajectory from that value. Since it arrives at the same place at time zero no matter what value x is picked at time $-T$, the value returned by the algorithm at time zero is the tail end of a path that has run for an infinitely long time and is therefore a "perfect" draw from π.

Typically, perfect sampling is carried out using *stochastically dominating processes* that bound or "sandwich" the process of interest for all time as depicted in Figure 6. At time $-T$, these bounding processes are at the extremes of the state space. At time $-b$, they have coupled together to form a single path that is followed forward to time 0. If such coupling is not achieved before time 0, T is increased and the bounding processes are resimulated using the same random inputs whenever they pass through previously visited time points. Note that when the bounding processes couple at some time $-b$, all possible sample paths of the process of interest have necessarily coupled at some earlier time $-a$. Note further that a path started anywhere in the state space at time $-T$ will be at the same place at time 0. This value at time 0 is a perfect draw from the stationary distribution of the Markov chain. Time $-T$ is known as a *backward coupling time*.

Forward coupling of sample paths does not yield the same result for reasons that we will not go into here. For more information on perfect sampling algorithms, we refer the reader to [5].

5.2 ε-Perfect Simulation for the Kac Model

We say that an algorithm is "ε-perfect" if the distance between the bounding processes depicted in Figure 6 is monotonically decreasing and if this distance is less than ε at time 0. In this Section we will describe an ε-perfect algorithm for sampling from $f(v, \infty)$ that is interesting mainly as a proof of concept for the approach that we hope will eventually be adapted to and useful for more interesting models.

Figure 6: Depiction of stochastically dominating (bounding) processes in perfect simulation.

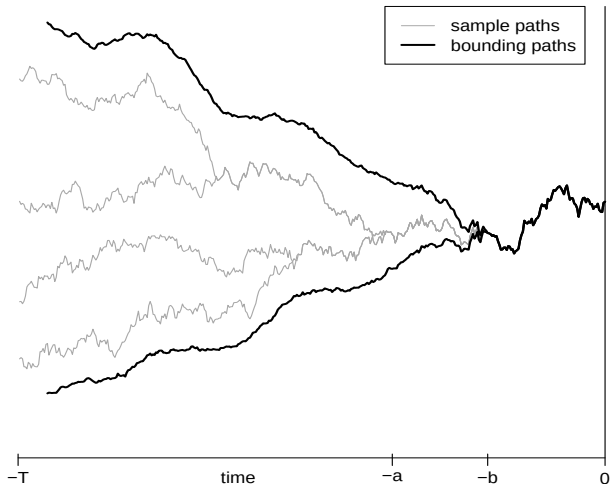

Consider, for the purpose of illustration, the case where we have only $N = 3$ particles. Simulating iid velocities v_1, v_2, and v_3 from $f_0(v)$ gives a kinetic energy $E = v_1^2 + v_2^2 + v_3^2$. As we propagate these 3 velocities forward in time, we create a realization of a 3-dimensional Markov chain that lives on the surface of a 3-dimensional sphere with radius $\sqrt{E}$. For the simple model considered in this paper, we can, without loss of generality, consider the velocity vector as a random walk on the sphere restricted to the first octant where $v_1, v_2, v_3 > 0$ and at the end of our simulation attach a negative sign to each velocity independently of the others with probability $1/2$.

For the purpose of achieving a coupling of sample paths walking on the sphere, we note that $v_i \cos\theta + v_j \sin\theta$ can be written as

$$v_i \cos\theta + v_j \sin\theta = A \sin(\theta + \varphi)$$

for some A and φ. In the case where $v_j = 0$, we may take $A = v_i$ and $\varphi = \pi/2$. In the case that $v_j \neq 0$, we may take $A = \sqrt{v_i^2 + v_j^2}$ and $\varphi = \tan^{-1}(v_i/v_j)$. Suppose that the current position of the random walk is (v_1, v_2, v_3) and that the particles randomly selected for updating are indexed by $(i, j) = (1, 2)$. Note then that $A = \sqrt{E - v_3^2}$ and, since v_2 is non-zero with probability 1, the new velocity for particle 1 is given by

$$v_1' = \sqrt{E - v_3^2}\ \sin(\theta + \varphi) \tag{12}$$

where $\theta \sim \text{unif}(0, 2\pi)$ and $\varphi = \tan^{-1}(v_1/v_2)$. Since $\theta + \varphi$ is uniformly distributed over an interval of length 2π, the distribution of v_1' in (12) is the same as that of $\sqrt{E - v_3^2}\ \sin(\theta)$.

Figure 7: Visualization of sample paths of the "Perfect Kac Walk" on a sphere in the first octant.

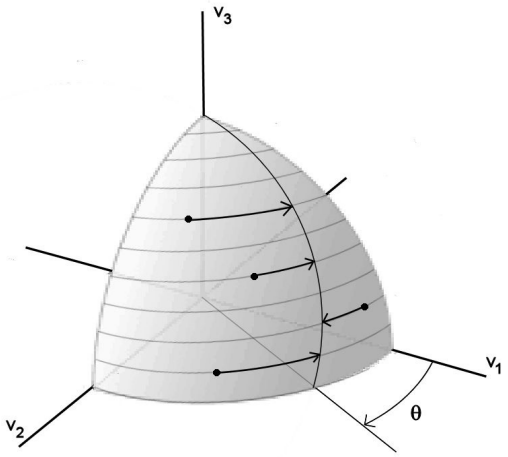

Thus, we will update v_1 as

$$v_1' = \sqrt{E - v_3^2}\ \sin(\theta) \tag{13}$$

where $\theta \sim \text{unif}(0, 2\pi)$ and then v_2 by setting

$$v_2' = \sqrt{E - v_3^2 - (v_1')^2}. \tag{14}$$

Four different values for (v_1, v_2, v_3) are visualized as points on the sphere in the first octant in Figure 7, as well as an angle θ drawn uniformly over $(0, 2\pi)$. All four points, **and indeed all points on the entire surface**, will map onto the depicted arc.

Suppose now that the next particles randomly selected for updating are indexed by $(i, j) = (3, 1)$. Given a new angle θ drawn uniformly over $(0, 2\pi)$, the points on the arc from Figure 7 will map to the smaller arc segment depicted in Figure 8.

It is easy to see that if the same indices (in the same or different order) are selected twice in a row, the arc will maintain its length and move "laterally". However, if at least one index is different from those in the previous step, it is easy to verify that the arc segment will, necessarily, shrink in length. In other words, we will be mapping all points on the sphere closer and closer together. Thus, the algorithm is closing in on a single point on the surface. To formalize the ε-perfect sampling algorithm for general N, we will follow the N corner points $\vec{c}_1, \vec{c}_2, \ldots, \vec{c}_N$, where $\vec{c}_i = \sqrt{E}\,\vec{e}_i$ and $\vec{e}_i$ is the ith standard basis vector for $\mathbb{R}^N$.

ε-Perfect Algorithm

Set an error tolerance $\varepsilon > 0$ and set $n = 1$.

1. Simulate and store $\theta_{-n} \sim \text{unif}(0, \pi/2)$.

Figure 8: Arc from Figure 7 mapped to shorter arc.

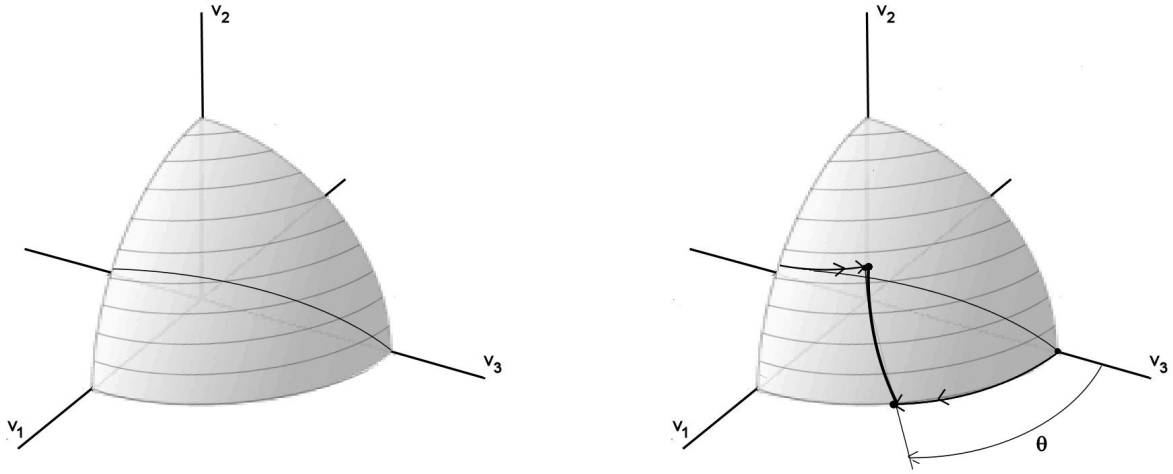

(a) Previous arc with rotated axes

(b) New angle and arc mapped to smaller arc

Randomly select a pair of distinct indices from $\{1, 2, \ldots, N\}$ and store as a 2×1 vector $\vec{P}_{-n}$ with entries denoted as $\vec{P}_{-n}(1)$ and $\vec{P}_{-n}(2)$. (Order is important. For example $\vec{P}_{-n} = (3, 5)$ is different from $\vec{P}_{-n} = (5, 3)$.)

Set $\vec{c}_i = \sqrt{E}\,\vec{e}_i$ for $i = 1, 2, \ldots, N$.

Let $t = n$.

2. Let $\vec{c}_i(j)$ denote the jth element of $\vec{c}_i$.

   For $i = 1, 2, \ldots, N$, let $e = [c_i(P_{-t}(1))]^2 + [c_i(P_{-t}(2))]^2$ and set

   $$\vec{c}_i(P_{-t}(1)) = \sqrt{e}\ \sin\theta_{-t}$$

   and

   $$\vec{c}_i(P_{-t}(2)) = \sqrt{e - [\vec{c}_i(P_{-t}(1))]^2}$$

   If $t = 1$, go to Step 3. Otherwise, let $t = t - 1$ and return to the beginning of Step 2.

3. Compute the Euclidean distance between all pairs of corner points. Let D be the maximum of these distances. If $D < \varepsilon$, compute the average corner point

   $$\vec{c} := \frac{1}{N} \sum_{i=1}^{n} \vec{c}_i,$$

   multiply each coordinate of $\vec{c}$ by -1 with probability $1/2$ independent of all other coordinates, output the new $\vec{c}$, and exit the algorithm. Otherwise, set $n = n + 1$ and return to Step 1.

In practice, it is more efficient, in the search for a backward coupling time, to step back further than one step at a time. It is generally advised to let $t = 2^n$ in, what is in this case,

Figure 9: Histogram of $100,000$ draws of v_1 Using the ε-perfect algorithm with $N = 50$ and $\varepsilon = 10^{-6}$.

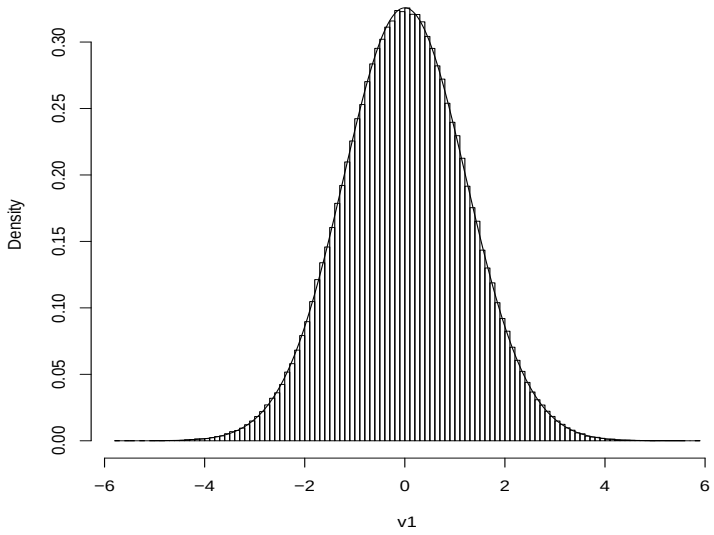

Step 1 of the ε-perfect algorithm. For simplicity, we use the Euclidean distance between corner points as a surrogate for arc length distance on the surface of the sphere since convergence to zero of the former implies convergence to zero of the latter.

In Figure 9 we show the results for $100,000$ draws of v_1 using the ε-perfect algorithm with $N = 50$ and $\varepsilon = 10^{-6}$ along with the curve given by (9). For this simulation run, the mean backward coupling time was 948.2 with a minimum of 422 and a maximum of 1811.

6 Conclusions

In this paper we leveraged fundamental properties of the Poisson process to simulate values drawn from the probability density function that is the solution to the Kac equation in an efficient way that avoids the tuning parameter Δt that is necessary for the Nanbu and Nanbu-Babovsky algorithms. Since our algorithm takes advantage of the fact that the particle 1 velocity is unchanged after a point in $(0, t)$, it expects to process a slightly lower number of collisions than Bird's DSMC algorithm. Furthermore, it appears to be performing better at capturing the tails of the distribution than the DSMC algorithm. This is likely due to the fact that we have a variable Poisson number of collisions in any time interval as opposed to the constant expected value.

We also introduced, to our knowledge, the first perfect sampling type algorithm to be applied to the Kac model as a proof of concept that has the potential to be expanded

to more interesting applications that are currently simulated using variants of the DSMC algorithm. A truly perfect ($\varepsilon = 0$) algorithm using a minorization approach (as in [6]) is possible but the increase in accuracy is not worth added complication, especially in light of the fact that we can set ε in line with the maximum attainable computer dependent precision.

References

[1] H. Babovsky and R. Illner. A convergence proof for Nanbu's simulation method for the full Boltzmann equation. *SIAM J. Numer. Anal.*, 26:45–65, 1989.

[2] G.A. Bird. *Molecular Gas Dynamics*. Clarendon Press, Oxford, 1994.

[3] E. Carlen, M. Carvalho, and M. Loss. Kinetic theory and the Kac master equation. *Entropy & the Quantum II, Contemporary Mathematics*, 552:1–20, 2011.

[4] E.A. Carlen, M.C. Carvalho, J. Le Roux, M. Loss, and C. Villani. Entropy and chaos in the kac model. *Kinetic and Related Models*, 3(1):85–122, 2010.

[5] G. Casella, M. Lavine, and C.P. Robert. Explaining the perfect sampler. *The American Statistician*, 55(4):299–305, 2001.

[6] J.N. Corcoran and R.L. Tweedie. Perfect sampling of ergodic Harris chains. *Ann. App. Prob.*, 11(2):438–451, 2001.

[7] M. Kac. Foundations of kinetic theory. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 3: Contributions to Astronomy and Physics*, pages 171–197. University of California Press, Berkeley, California, 1956.

[8] M. Krook and T.T. Wu. Formation of Maxwellian tails. *Physical Review Letters*, 36(19):1107–1109, 1976.

[9] K. Nanbu. Direct simulation scheme derived from the Boltzmann equation. i. monocomponent gases. *Journal of the Physical Society of Japan*, 49(5):2042–2049, 1980.

[10] W. Wagner. A convergence proof for Bird's direct simulation Monte Carlo method for the Boltzmann equation. *J. Statist. Phys.*, 66:1011–1044, 1992.